



## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification <sup>6</sup> :	A2	(11) International Publication Number: <b>WO 97/41674</b>
H04L 29/06		(43) International Publication Date: 6 November 1997 (06.11.97)
(21) International Application Number:	PCT/US97/07271	(81) Designated States: AU, CA, GB, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
(22) International Filing Date:	29 April 1997 (29.04.97)	
(30) Priority Data:		<b>Published</b>
08/641,399	30 April 1996 (30.04.96)	<i>Without international search report and to be republished upon receipt of that report.</i>
(71) Applicant:	3COM CORPORATION [US/US]; 5400 Bayfront Plaza, Santa Clara, CA 95052-8145 (US).	
(72) Inventors:	BINDER, James, Stuart; 1277 Elkwood Drive, Milpitas, CA 95035 (US). SIDENBLAD, Paul, James; 10190 Stonydale, Cupertino, CA 95014 (US). BAKER, Roman, Gabriel; 339 Kenmore Avenue, Sunnyvale, CA 94086 (US). CONNERY, Glenn, William; 655 S. Fair Oaks #B301, Sunnyvale, CA 94086 (US).	
(74) Agents:	LEBLANC, Stephen, J. et al.; Townsend and Townsend and Crew L.L.P., 8th floor, Two Embarcadero Center, San Francisco, CA 94111 (US).	

(54) Title: PACKET FILTERING BASED ON SOCKET OR APPLICATION IDENTIFICATION

## (57) Abstract

A network adaptor driver includes an algorithm for determining special handling of packets by examining socket or port indications for the packets and a means for setting a special handling field in the header of outgoing packets. Specific embodiments include a table for connections that include packets without socket indications, a transmit FIFO list for determining whether a transmit FIFO is full, and a low priority receive queue. A computer readable media allowing a computer to practice the method of the invention is disclosed.

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

PACKET FILTERING BASED ON  
SOCKET OR APPLICATION IDENTIFICATION

5

APPENDIX

This application is being filed with a paper appendix consisting of 95 pages containing a source code listing for an embodiment of the invention. A microfiche reproduction of this appendix will be submitted upon notice of allowance.

BACKGROUND OF THE INVENTION

This application is a Continuation-In-Part of co-assigned pending U.S. Application Serial No. 08/313,674, entitled METHOD AND APPARATUS FOR CONTROLLING LATENCY AND JITTER IN A LOCAL AREA NETWORK WHICH USES A CSMA/CD PROTOCOL.

The current invention relates to the field of electronic circuits. More particularly, the current invention relates to transmission of information between digital devices over a communications medium.

The present invention relates to the field of data communication over a computer network. A very wide variety of types of computer networks exist, each having many variations in particular implementations. The present invention will be described with reference to a particular types of computer networks, operating systems, and network interfaces. This should not be taken to limit the invention, and it will be apparent to those of skill in the art that the invention has applications in many different types of computer networks. The invention therefore should not be seen as limited except as provided in the attached claims.

Digital computer networks have become ubiquitous in academic, industry, and office environments. A number of

different aspects of computer networks are discussed in co-assigned pending U.S. applications serial nos. 08/313,674; 08/542,157; 08/506,533; and 08/329,714 each of which are incorporated herein by reference.

5           Most of the commonly used networks today operate in accordance with a layered protocol suite. In a layered protocol suite, various network tasks are divided among different pieces of hardware and software in order to allow maximum flexibility and reliability between software units.

10           Fig. 1 shows a simplified block diagram of a layered protocol suite as might exist between a personal computer 10 and a server computer 20 communicating via a network 15, such as two Windows-compatible type computers in an office environment running over a LAN. In a simplified form, such a 15 system can be thought of as comprising four semi-independent layers running on each computer.

20           A top layer, layer 4, is the application layer. The application layer generally is occupied by software running on the central processor of computer 10. This software might include video conferencing software, e-mail software, the file system, or any other software that accesses network 15. The application layer 4 on computer 10 communicates data, such as a video conferencing signal, with an application layer on computer 20, by passing that data down through the protocol stack on computer 10 to the lowest layer, layer 1. Layer 1 provides a physical connection to the network and sends the data to the layer 1 device on computer 20, which then passes the data back up to its application layer.

25           Application layer 4 communicates data down to protocol layer 3. Protocol layer 3 includes software necessary to perform high level network functions such as accepting a data file from an application layer, dividing that file into packets, and attaching to those packets source and destination address information so that packets can be effectively transmitted over the network. Some examples of 35 different protocol layer software are TCP/IP or IPX.

The protocol layer communicates with an adapter driver layer 2. Layer 2 generally includes software for interfacing between adapter 1 and the protocol layer 3. In prior art layered implementations, layer 2 is generally low level software whose primary purpose is to act as the interface between the protocol layer and the particular adapter hardware that is connected to or a part of computer system 10. In prior art networks, layer 2 software does not examine the contents or address of any data passing through it, but merely formats that data for the particular adapter to place on the network. Some examples of driver layer software are NDIS2 or ODI.

Adapter 1 is generally thought of as a piece of hardware for communicating signals over the network media. Adapter 1 will have different configurations depending on whether the network media is a co-axial cable, twisted-pair wire, fiber optic cable, or radio waves. The adapter may be a plug in board that connects to computer 10's system bus. Once adapter 1 receives a packet of data to be transmitted over network 15, it sends a packet over network 15 which delivers it to an adapter 1A in addressed computer 20. Once data reaches address computer 20, adapter 1A passes the packet up the layered stack first to driver layer 2A, then to protocol layer 3A, which then passes to application layer 4A.

An adapter generally includes circuitry and connectors for communication over a network medium and translates data from the digital form used by the computer circuitry a form that may be transmitted over a medium, such as a waveform. A driver is a set of instructions resident on a device that allows the device to accomplish various tasks as defined by different network protocols. Drivers are generally software programs stored on the computer in a manner that allows the drivers to be modified without modifying other hardware.

In a layered network system such as shown in Fig. 1, various networking tasks are tightly defined and performed at various network layers. In prior art systems, driver 2, for

example, receives packets of data and forwards them immediately to adapter 1 without examining the contents of the packets of data and without ever changing any of the contents of the packets of data. Driver 2 essentially exists to be an interface between a wide variety of different hardware equipment that may be provided as adapter 1 and the limited number of different network protocols 3 that might be supported by computer system 10. In prior art systems, driver 1 simply sends each packet as it receives it or sends packets in a strict FIFO order where the first packet received from protocol 3 is the first packet sent out to adapter 1 and over the network.

The layered system shown in Fig. 1 has worked well in network environments where a majority of network traffic consists of text or computer instructions which do not have particular time-critical requirements and where a computer, such as 10, only allowed one application to access the network at a time.

However, two developments in personal computer technology have shown that this system has certain limitations. The first is the wide-spread use of socket capable protocols and operating systems in networked environments. In a socket system, multiple applications running simultaneously on computer 10 can simultaneously send and receive data on network 15. Protocol layer 3 keeps track of for which application a particular data packet is destined by assigning different socket numbers. Socket numbers are requested by applications from protocol layer 3. All data packets sent and received include this number in their headers and thus can be delivered to the right application. The location of the socket identification within the packet is determined by the packet's network protocol.

The second development that has highlighted the limitations of prior art adapter drivers is that computer network traffic has increasingly begun to include real-time data such as digitized audio signals and digitized video signals, such as signals that might be used for a real time

video conferencing system. In a sockets compatible computer system, this real-time network data might be attempting to use the network at the same time that other, high volume but non-time critical, computer data is being placed on the network by 5 a different application. Without a way for the computer system to ensure that real time video and audio data is delivered at a high priority while delays are allowed to happen in data streams that are not time critical, the real-time data can be delayed unacceptably long, hurting multi-  
10 media performance.

What is needed is a method and apparatus allowing a computer system such as 10 to send real-time data on a network at a higher priority than non-time-critical data and that requires minimal modification to other network hardware or  
15 software.

#### SUMMARY OF THE INVENTION

The present invention is an adapter driver for use with a network adapter card designed to recognize an identification field from an application layer program and to 20 assign special handling or filtering to packets received from said application program based on the identification code. In a specific case, this filtering includes assigning a priority to individual packets. According to a specific embodiment, a driver assigns one of two priorities: a high priority (HP) or  
25 a low priority (LP). With the assignment of priority, a driver according to the present invention can manage and determine in what order packets are passed to an adapter 1 such that high priority packets are transmitted on the network before low priority packets are transmitted on the network.

According to a further embodiment of the invention, a driver according to the invention may also change data in a packet's header, thereby identifying special handling indications of said packet to other communications devices in  
30 a network.

According to a further embodiment of the invention, a driver according to the invention may detect a high priority  
35

bit in a received packet's header and thereby deliver that packet to the layer 3 protocol ahead of other received packets.

According to a specific embodiment, a driver is designed to operate in a computer system which can function with a number of different network protocols. The driver includes instructions allowing it to read a network protocol identifier from a packet of data and then to determine the source and destination address from said packet based on the protocol used to encode said packet. An example of one type of system in which the invention may be used is the Windows(TM) operating system, developed and sold by Microsoft(TM) Corporation. In the Windows operating system, applications communicating data to a network identify packets of data with a socket identification. The location of this socket identification within the packet is determined by the packet's network protocol. (In other operating systems, the name for what Window's calls a socket is a port.)

In a Windows embodiment, a driver according to the invention reads a packet protocol and from the packet protocol determines the location of a socket identifier. The driver then reads the socket identifier and from it determines the priority of a packet. According to a specific embodiment, packet priority may be based on assigning particular ranges of socket identifiers to particular priorities. Such assignments may be stored in a configuration file on a computer such as computer 10 to be read by the adapter driver at startup.

A further example of an environment in which the invention may be employed is the Apple operating system. In the Apple operating system, applications also assign an identifier to every packet of data that is sent to the protocol layer.

An adapter driver according to one embodiment of the invention may include one or more priority transmit queues. These queues, according to one embodiment, are established by the driver in computer 10's main system memory. According to

a different embodiment, these queues may be placed in a dedicated memory on the adapter card. According to still another embodiment, a driver may handle high priority packets by interacting directly with memory on the adaptor card and without using priority queues. Where a driver interacts with just one priority queue, the driver includes control means allowing it to place packets at any location in the queues depending on the packet's priority. Where a driver interacts with multiple queues, a packet is located within a queue indicating that packet's priority.

According to one embodiment of the invention, an adapter driver according to the invention may be implemented entirely in software and may be designed to run on existing computer systems with existing higher layer protocol drivers and existing adapters. In this way, the advantages of the invention may be achieved by modifying only the adapter driver software on a computer such as 10, which is a relatively inexpensive investment to gain the performance benefits of the invention. This software may be embodied in computer readable executable instructions written onto a fixed medium such as a disk, for distribution to individual computer systems either via removable or fixed disks or over a transmission medium.

In a different embodiment, the invention may be used in conjunction with other modified network software and hardware to achieve further enhancements in performance. These other modified network elements could include network switching devices that recognize a priority field set by an adapter driver according to the invention and that provide special handling for high priority packets.

In a further embodiment, a driver according to the invention is optimized to run on a modified adapter card which is designed to recognize and provide special handling for high priority packets.

Finally, the essential technique of the invention includes using an adaptor driver to examine higher layer information in a packet in order to mark the packet for

various special handling. In addition to priority handling, this special handling could include encryption, compression, or security blocking. This special handling is distinct from prior art handling of network data because it is done on a 5 packet by packet basis by the adaptor driver layer of the protocol. The present invention is intended to encompass these variations.

#### BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram illustrating a layered network 10 protocol of a type in which the present invention may be effectively employed.

Fig. 2 is a block diagram of a computer system and an adapter with a driver according to an embodiment of the present invention.

15 Fig. 3 is a diagram of an example of a packet format which may be effectively handled by the present invention.

Fig. 4 is a flow chart of a method of the driver for special handling of the packets according to the present invention.

20 Fig. 5 is a flow chart of a method by which a driver according to the present invention sends packets to an adapter.

25 Fig. 6 is a flow chart of a method by which a driver according to the present invention receives and buffers low priority packets.

Fig. 7 is a block diagram of a computer system with an ability to read a computer readable medium allowing the system to operate in accordance with the invention.

30 Fig. 8 shows a system block diagram of computer system 10 used to execute the software of the present invention.

#### DESCRIPTION OF THE PREFERRED EMBODIMENT

Fig. 2 illustrates a computer system 10 containing

an adapter 1, an adapter driver 2 according to an embodiment the invention, a protocol layer 3, and application layer 4. Driver 2 has associated with it a protocol detection algorithm (PDA) 50, a high priority transmit queue (HPTXQ) 55, and a low priority transmit queue (LPTXQ) 60. According to one embodiment, driver 2 also has associated with it an adapter FIFO transmit list 65 and a low priority receive queue (LPRXQ) 70.

Adapter 1 has associated with it an adapter transmit FIFO 110 and an adapter receive FIFO 115. Adapter 1 communicates packets through driver 2 to computer system 10 and transmits and receives packets from network 15.

Fig. 3 illustrates an example of a network protocol packet which may be effectively handled by an adapter driver according to the present invention. Fig. 3 illustrates a packet 200, containing packet header 202. Packet header 202 contains, a protocol identifier 210, a source socket identifier 204, a source address 206, a destination socket identifier 208, and a destination address 210. According to the present invention, packet header 202 may also contain a locally administered bit 212 which is used by an embodiment of the present invention to indicate special handling of the packet in the network 15.

The present invention according to one embodiment may also determine the sockets of packets where not every packet includes a socket identification, such as UDP protocol packets. In the case, a driver 2 according to the invention, maintains a table of UDP identifiers along with the socket to which a particular UDP connection is assigned.

Fig. 4 illustrates a flow chart of the operation of adapter driver 2 illustrating the method of the special handling of packets based on their source socket or their destination socket address according to an embodiment of the present invention. A packet is received by driver 2 from protocol 3 (Step S2). Driver 2, according to the invention, examines each packet it receives on transmission from protocol

3 to determine special handling needed by the packet. First, driver 1 looks for a protocol identifier in every packet it receives so as to determine which protocol created the packet (Step S4). In one embodiment this is done by a subroutine 5 within driver 2 referred to as the priority determining algorithm (PDA) 50. Examples of possible packet protocols are PCP, EDP, UDP, or IPX.

Once PDA 50 has determined which protocol formatted a packet, PDA 50 then knows how and from where to read the 10 header of the packet and to determine the source and destination socket numbers for the packet (Step S6).

After determining the source socket number for a packet, PDA 50 determines what special handling instructions should be assigned to the packet. This is determined by PDA 15 50 reading from a configuration memory 52 the ranges of socket numbers that are assigned special handling instructions and determining in which range the socket number of the current packet falls (Step S8).

According to an embodiment of the current invention, 20 driver 2 maintains at least two transmit queues of different priorities in computer system 10 such as high priority transmit queue 55 and low priority transmit queue 60. Based on the special handling determination made in Step S8 (Step S10), driver 2 places the packet in one of these two transmit 25 queues. In the current specific example, driver 2 will place a packet in either LPTXQ 60 (Step S12) or HPTXQ 55 (Step S16), depending on the priority determination made for the packet. According to a further embodiment, driver 2 will add priority 30 data to the packet header (Step S14) before the packet is placed in one of the two queues. This priority data may be added to either all high priority packets, all low priority packets, or all packets.

Figure 5 illustrates a further method of the 35 invention whereby driver 2 also determines which packets will be sent to adapter 1 for transmission on network 15. According to one example, this determination is simple and is

based on whether any packets are present in HPTXQ 55. If any packets are present HPTXQ 55, driver 2 sends those packets to adapter 1 immediately (Step T14), and allows any packets within LPTXQ 60 to wait until there are no more HP packets in queue 55 remaining to be transmitted.

According to a related further embodiment of the invention, driver 2 ensures that at least some LP packets are transmitted even when there is a very high volume of HP packets being received by driver 2, referred to herein as fairness (Step T12). In this embodiment, driver 2 maintains a count MPC of HP packets transmitted without interruption, and when that count reaches a certain value, driver 2 sends out a packet from LPTXQ 60 to adapter 1 if one is present.

According to a further embodiment of the invention, driver 2 is enabled to effectively schedule HP packets even when an adapter such as 1 includes a large transmit FIFO 110. In order to operate effectively with prior art adapters 1, which are not aware of the scheduling of HP and LP packets, driver 2 takes responsibility of halting LP packets being transmitted to FIFO 110 when that FIFO becomes too full. Driver 2 does this to avoid the situation where FIFO 110 contains a large number of LP packets and driver 2 receives an HP priority packet to transmit. When used with prior art adapters, driver 2 has no way to force the HP packet ahead of other packets stored in FIFO 110, and the HP packet would have to wait.

A driver 2 according to this embodiment of the invention, avoids this problem by maintaining an adapter FIFO transmit list 65, in which driver 2 logs the presence of all packets sent to adapter 110. Driver 2 then examines the FIFO free space value returned from adapter 110 to determine which packets that it has stored have been sent by the adapter onto network 15. When driver 2 determines that a certain number of low priority packets reside on transmit FIFO 110 and have not been sent, driver 2 halts forwarding any further low priority packets to adapter 1 until the number of low priority packets stored in transmit FIFO 110 again falls below a certain

threshold (Step T16).

In a further embodiment of the invention, driver 2 first checks the packet traffic to determine if mixed HP and LP packets are being sent, before it halts transmission of 5 packets to FIFO 110. If mixed traffic is not being sent, driver 2 suspends this technique and sends packets normally until transmit FIFO 110 is full.

Figure 5 is a flow chart of the method by which an 10 adapter including each of these alternative embodiments sends packets from its HPTXQ and LPTXQ to adapter 1.

Figure 6 illustrates a method of driver 2 for receiving packets according to a further embodiment of the invention. In a system such as 10 on a network, often several 15 packets will be received in rapid succession by adapter 1 and will be present for service together in RX FIFO 115. According to this embodiment, a driver 2 examines the header of each incoming packet when multiple packets are received, looking for a priority indication (Step U4). Because of the 20 operation of RX FIFO 115, driver has no choice but to receive the packets out of FIFO 115 in FIFO order and cannot receive HP packets first. Driver 2 therefore stores any LP packet it receives in LPRXQ 70 (Step U10) and continues receiving 25 packets from FIFO 115. When an HP packet is received, driver 2 immediately passes that packet up to protocol layer 3 (Step U8). When all packets in FIFO 115 have been received by 30 driver 2, driver 2 then passes any LP packets stored in LPRXQ 70 to protocol layer 3 (Step U8). According to one embodiment, driver 2 first determines whether mixed priority packets are being received over network 15 (Step U6) before instituting storing packets in LPRXQ 70.

According to one embodiment of the invention, driver 2 may be embodied in computer executable code on a fixed computer medium, either for delivery to a computer 10 as a fixed disk or for delivery over a network. When loader on an 35 appropriately configured computer, this code will cause the computer and associated hardware to perform in accordance with

the methods of the invention. It should be understood that many of the above described individual features of the invention are configurable and may be enabled or disabled by configuration codes supplied to driver 2 by computer system 5 10.

An alternative embodiment of the invention includes just one TXQ, such as 55 and includes control circuitry allowing driver 2 to place packets at any location within queue 55 depending on the packets priority, which higher 10 priority packets placed in the queue so that they will be transmitted first to adapter 1.

Fig. 7 illustrates an example of a computer system used to execute the software of the present invention. Fig. 7 shows a computer system 10 which includes a monitor 703, screen 705, cabinet 707, keyboard 709, and mouse 711. Mouse 111 may have one or more buttons such as mouse buttons 113. Cabinet 107 is shown housing a disk drive 715 for reading a CD-ROM or other type disk 117. Cabinet 707 also houses familiar computer components (not shown) such as a processor, 20 memory, disk drives, and the like, as well as an adaptor 1 for connection to a network medium 15.

Fig. 8 shows a system block diagram of computer system 10 used to execute the software of the present invention. As in Fig. 7, computer system 10 includes monitor 703 and keyboard 709. Computer system 10 further includes subsystems such as a central processor 722, system memory 724, I/O controller 726, display adapter 728, serial port 732, disk 736, network interface 738, adaptor 1 and speaker 740. Computer readable media such as memory, hard disks, floppies, 25 CD-ROMs, tapes, flash memory, and the like may be used to store a computer program including computer code that implements the present invention. Other computer systems suitable for use with the present invention may include additional or fewer subsystems. For example, another computer system could include more than one processor 722 (i.e., a multi-processor system) or a system may include a cache 30 memory. 35

Arrows such as 722 represent the system bus architecture of computer system 10. However, these arrows are illustrative of any interconnection scheme serving to link the subsystems. For example, speaker 740 could be connected to 5 the other subsystems through a port or have an internal direct connection to central processor 722. Computer system 10 shown in Fig. 8 is but an example of a computer system suitable for use with the present invention. Other configurations of subsystems suitable for use with the present invention will be 10 readily apparent to one of ordinary skill in the art.

The invention has now been explained with reference to specific embodiments. Other embodiments will be apparent to those of skill in the art. In particular, method steps have been grouped and labeled as being part of various 15 sub-methods in order to increase clarity of the disclosure, however, these steps could be differently grouped without changing the essential operation of the invention. It is therefore not intended that this invention be limited, except as indicated by the appended claims.

WHAT IS CLAIMED IS:

1           1. A network adapter driver comprising:  
2                 an interface for transmitting formatted packets of  
3                 data to and from a higher protocol layer;  
4                 a special handling detection algorithm for reading  
5                 the protocol identifier from a packet of data, for reading a  
6                 socket identifier from said packet of data, and for  
7                 determining a special handling indication for said packet  
8                 based on said socket identifier; and  
9                 a transmit interface for transmitting packets of  
10                data to a network adapter with special handling based on the  
11                special handling indication of said packets.

1           2. A network adapter driver according to claim 1  
2           further comprising:  
3                 configuration memory for storing special handling  
4                 indications based on ranges of source socket identifiers of  
5                 said packets.

1           3. A network adapter driver according to claim 1  
2           further comprising:  
3                 a packet header modification algorithm for modifying  
4                 a data field in a packet header in order to indicate a special  
5                 handling indication of said packet.

1           4. A network adapter driver according to claim 1  
2           further comprising:  
3                 a table for storing socket indications for protocols  
4                 that include packets that do not have socket indications in  
5                 each packet header.

1               5. A network adapter driver according to claim 1  
2       wherein said adapter driver is a software module that may be  
3       installed on a computer system with an adapter without any  
4       modifications to said computer system or said adapter.

1               6. A method for providing special handling of  
2       packets in an adapter driver comprising the steps of:

3               receiving a packet of data from a higher protocol  
4       layer, said packet containing a packet header;

5               reading from said packet header and indication of  
6       the protocol that constructed that packet;

7               using said protocol to determine where in said  
8       header a source or a destination socket identifier is located  
9       and reading said socket identifier(s);

10              using said identifier(s) to determine special  
11       handling for said packet; and

12              transmitting said packet to an adaptor in accordance  
13       with said special handling.

1               7. The method according to claim 6 further  
2       comprising:

3               adding data to said packet header indicating said  
4       special handling before transferring said packets to an  
5       adaptor.

1               8. A fixed computer readable medium containing  
2       computer executable program code that when loaded into an  
3       appropriately configured computer system will cause the  
4       computer to perform the method of claim 6.

1                 9. A fixed computer readable medium containing  
2 computer executable program code that when loaded into an  
3 appropriately configured computer system will cause the  
4 computer to perform the method of claim 7.

1                 10. A network adapter driver comprising:  
2                         an interface for transmitting formatted packets of  
3 data to and from a higher protocol layer;  
4                         a priority detection algorithm for reading the  
5 protocol identifier from a packet of data, for reading a  
6 socket identifier from said packet of data, and for  
7 determining a priority indication for said packet based on  
8 said socket identifier; and  
9                         a transmit interface for transmitting packets of  
10 data to a network adapter with special handling based on the  
11 priority indication of said packets.

1                 11. A network adapter driver according to claim 10  
2 further comprising:  
3                         a priority transmit queue interface for storing  
4 packets to be transmitted into locations into a plurality of  
5 transmit queues based on a priority indication of said  
6 packets.

1                 12. A network adapter driver according to claim 10  
2 further comprising a receive queue interface for temporarily  
3 storing received packets that have a low priority indication.

1                 13. A network adapter driver according to claim 10  
2 further comprising:  
3                         a transmit FIFO control interface for receiving  
4 signals indicating the amount of free space in an adapter

5       transmit FIFO; and

6                    a transmit FIFO list memory for storing indications  
7           of which packets in a transmit FIFO have been sent wherein  
8           said adapter driver halts placing additional low priority  
9           packets in said transmit FIFO when a specified number of  
10          packets are determined to be residing in said transmit FIFO.

1                   14. A network adapter driver according to claim 10  
2          further comprising:

3                   configuration memory for storing priority  
4          indications based on ranges of source socket identifiers of  
5          said packets.

1                   15. A network adapter driver according to claim 10  
2          further comprising:

3                   a packet header modification algorithm for modifying  
4          a data field in a packet header in order to indicate a  
5          priority of said packet.

1                   16. A network adapter driver according to claim 10  
2          further comprising:

3                   a table for storing socket indications for protocols  
4          that include packets that do not have socket indications in  
5          each packet header.

1                   17. A network adapter driver according to claim 10  
2          further comprising a first transmit queue interface for a high  
3          priority transmit queue and a said second transmit queue  
4          interface for a low priority transmit queue.

1               18. A network adapter driver according to claim 10  
2       wherein said adapter driver is a software module that may be  
3       installed on a computer system with an adapter without any  
4       modifications to said computer system or said adapter and  
5       wherein said queues are allocated by said driver from said  
6       computer system's system memory.

1               19. A method for scheduling transmission of packets  
2       in an adapter driver comprising the steps of:

3               receiving a packet of data from a higher protocol  
4       layer said packet containing a packet header;

5               reading from said packet header and indication of  
6       the protocol that constructed that packet;

7               using said protocol to determine where in said  
8       header a source socket identifier is located and reading said  
9       socket identifier;

10              using said identifier to determine a priority for  
11       said packet; and

12              placing said packet into a transmit queue in a  
13       location indicated by said priority.

1               20. The method according to claim 19 further  
2       comprising:

3               sending packets out of said higher priority queue  
4       locations before sending them out of lower priority locations.

1               21. The method according to claim 19 further  
2       comprising:

3               adding data to said packet header indicating said  
4       priority before placing said packets a queue.

1               22. The method according to claim 19 further  
2 comprising:

3               adding data to said packet header indicating said  
4 priority before placing said packets a queue.

1               23. A fixed computer readable medium containing  
2 computer executable program, which, when loaded into an  
3 appropriately configured computer system will cause the  
4 computer to perform the method of claim 19.

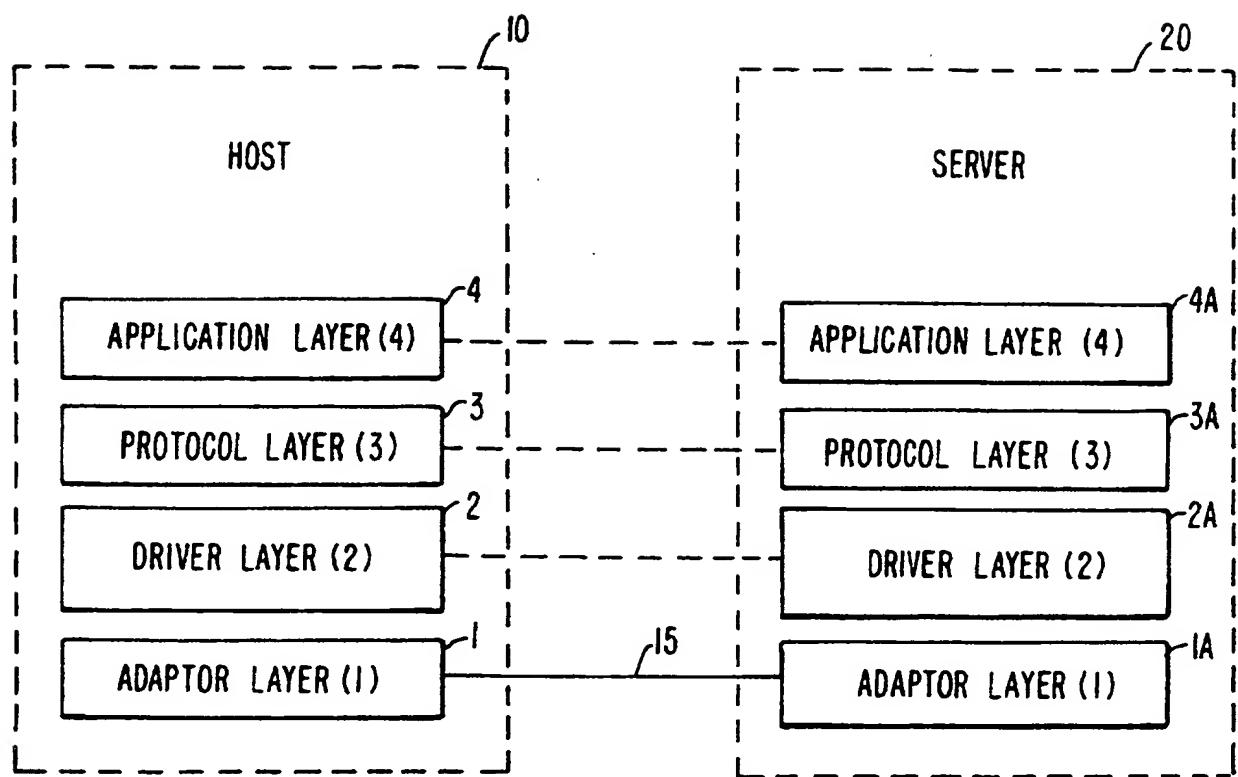


FIG. 1.

2/6

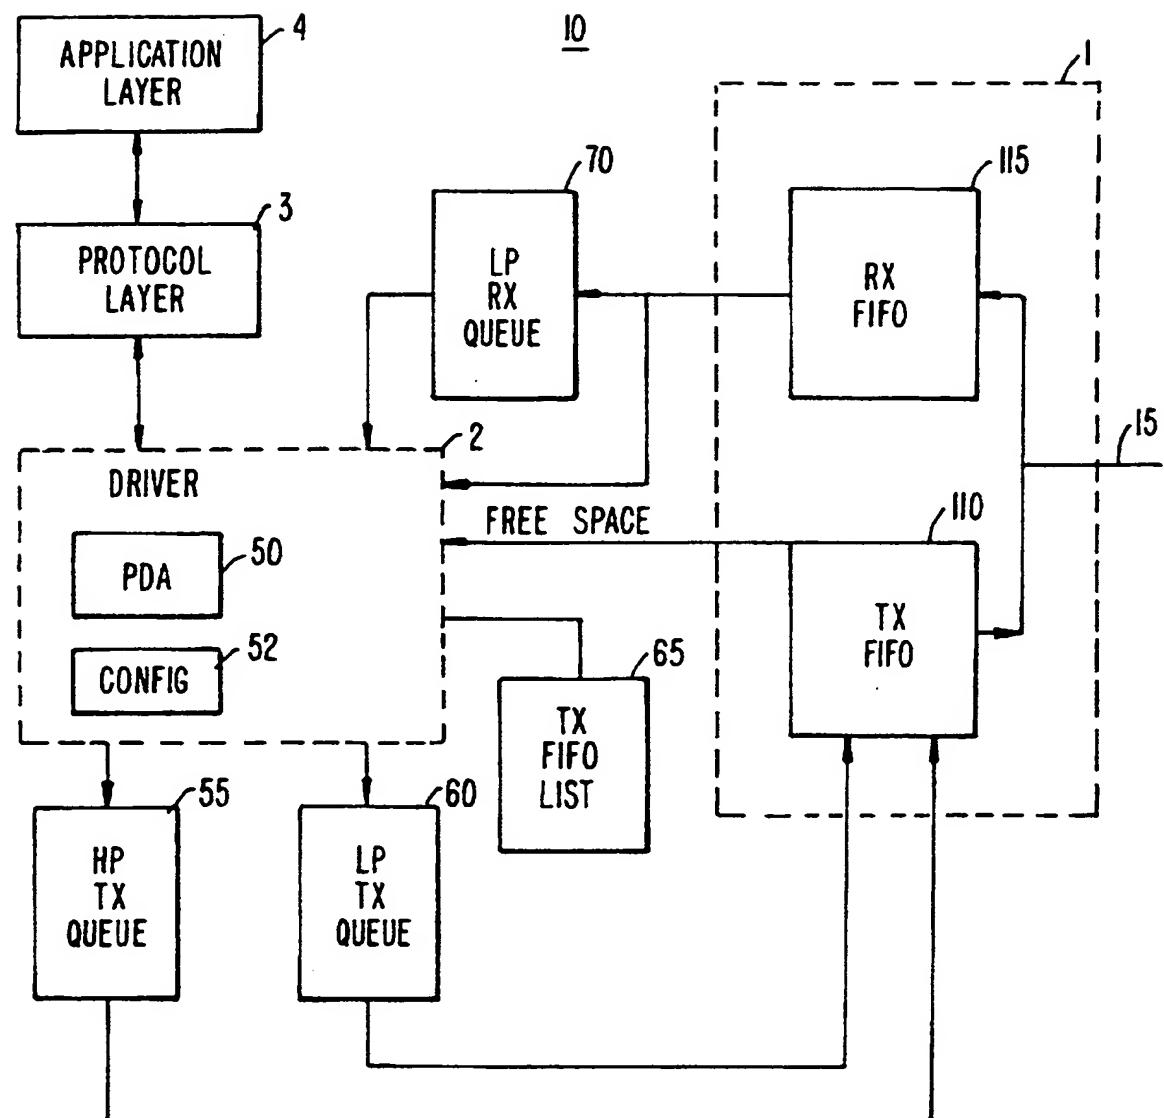


FIG. 2.

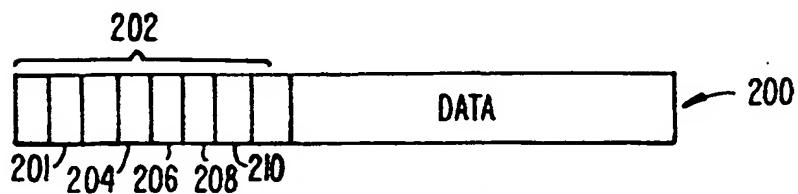


FIG. 3.

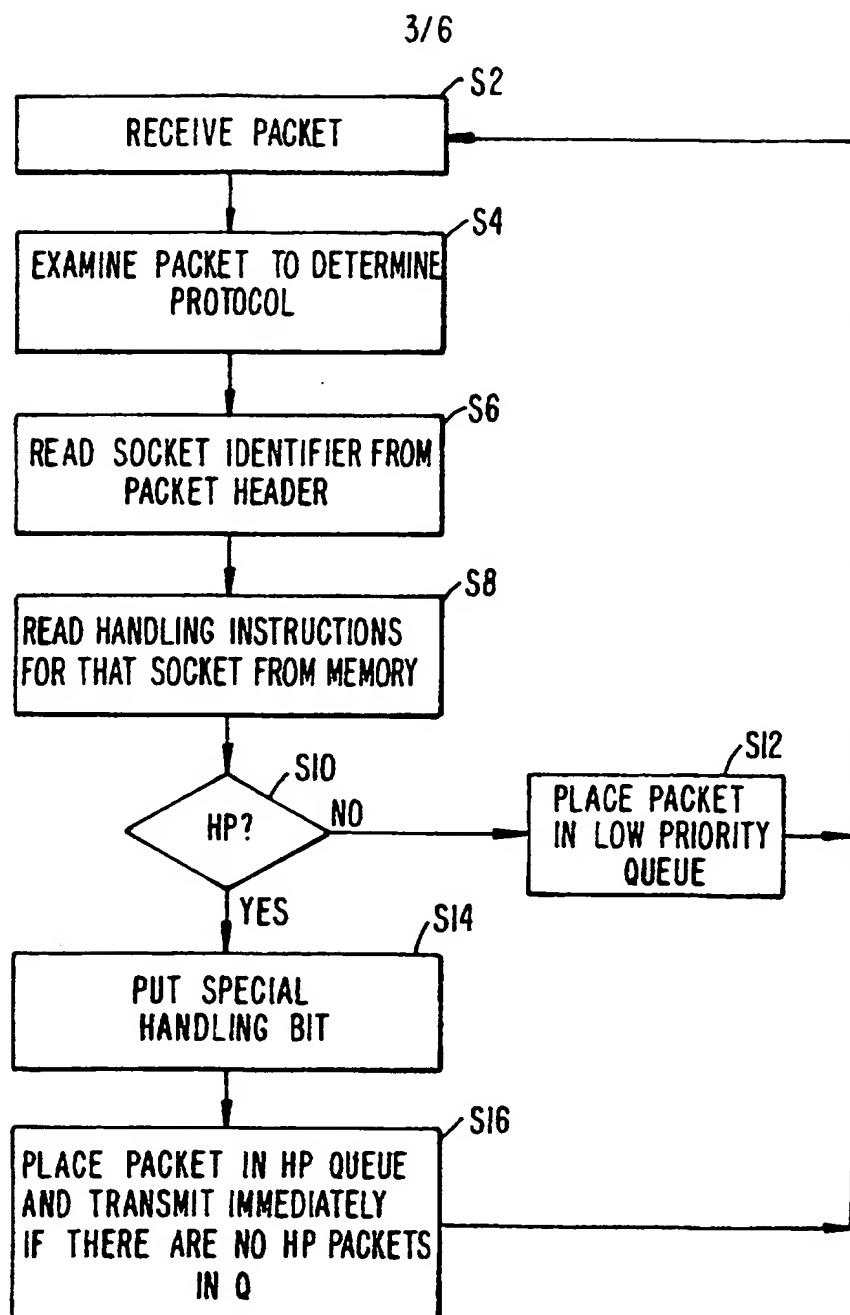


FIG. 4.

4/6

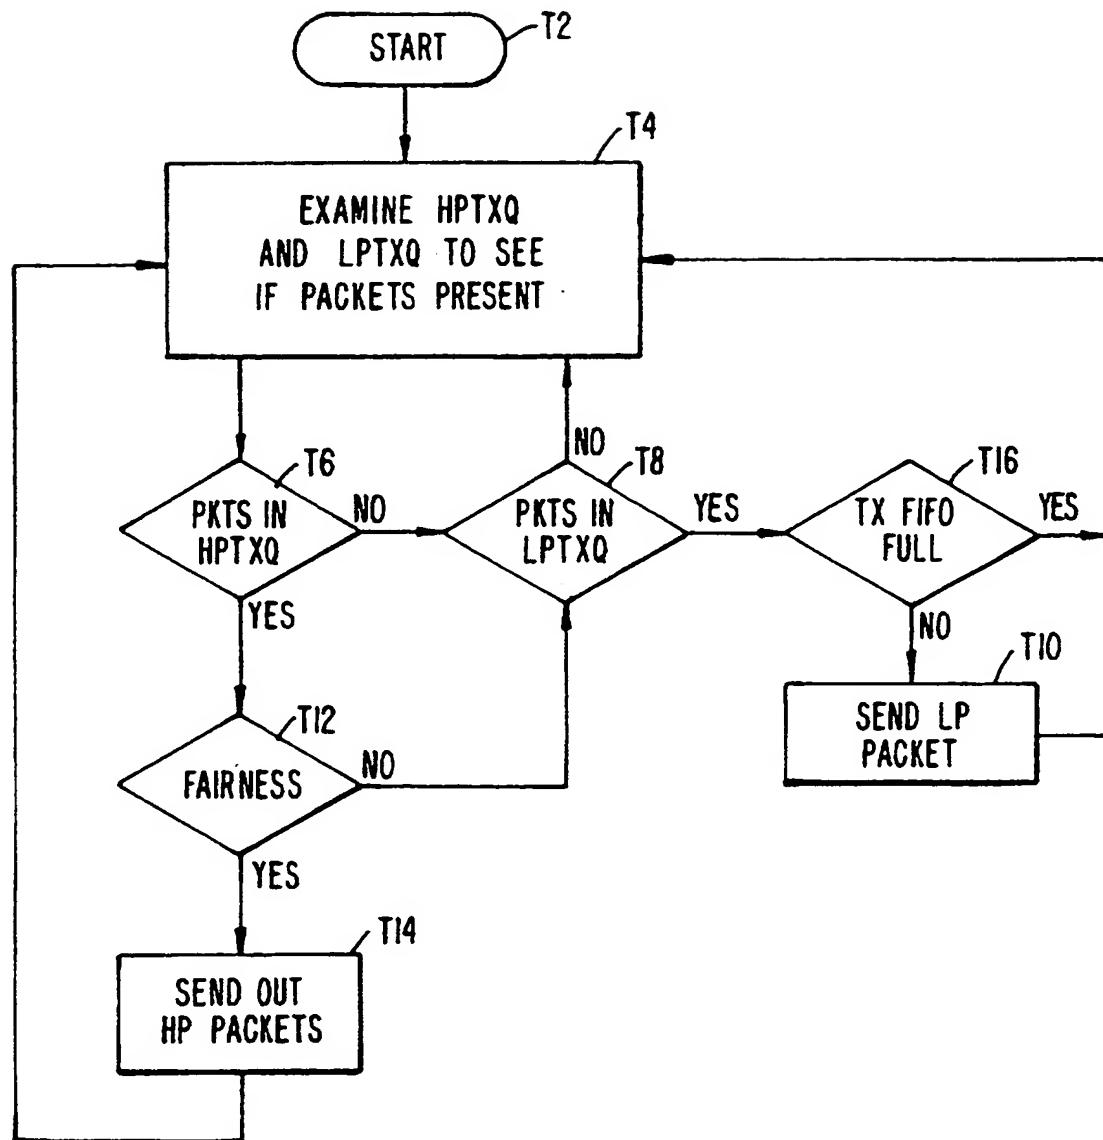


FIG. 5.

5/6

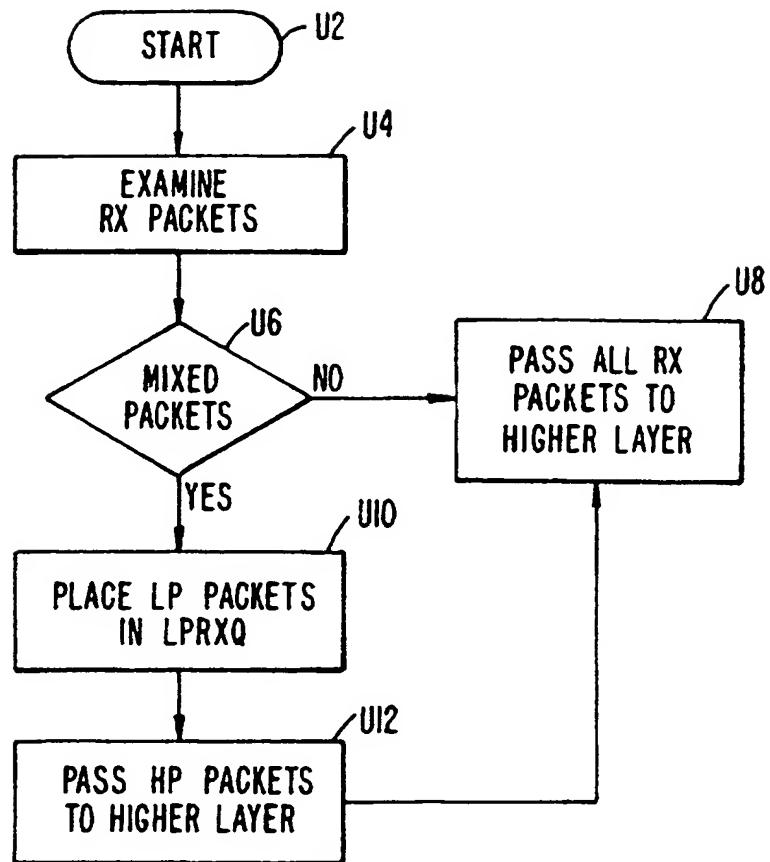
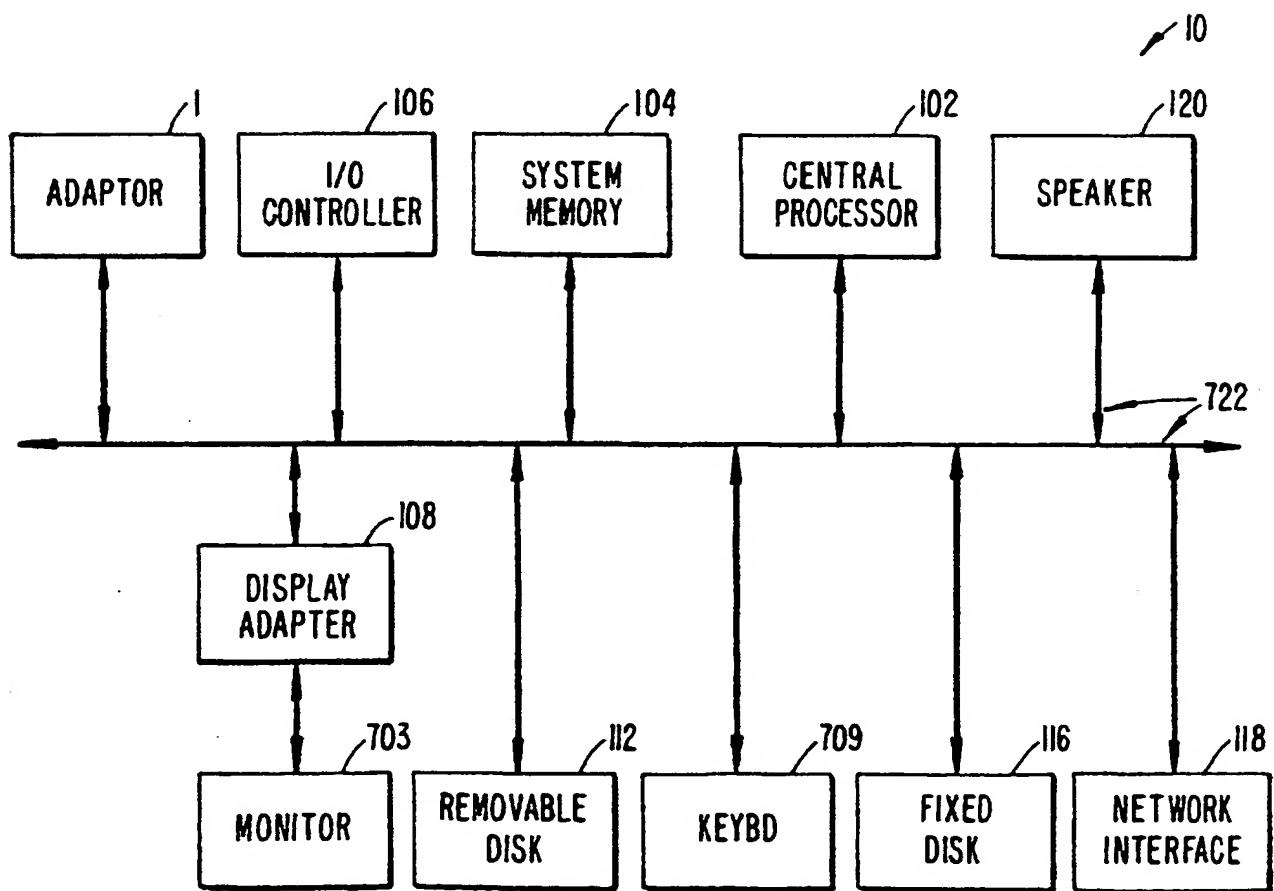
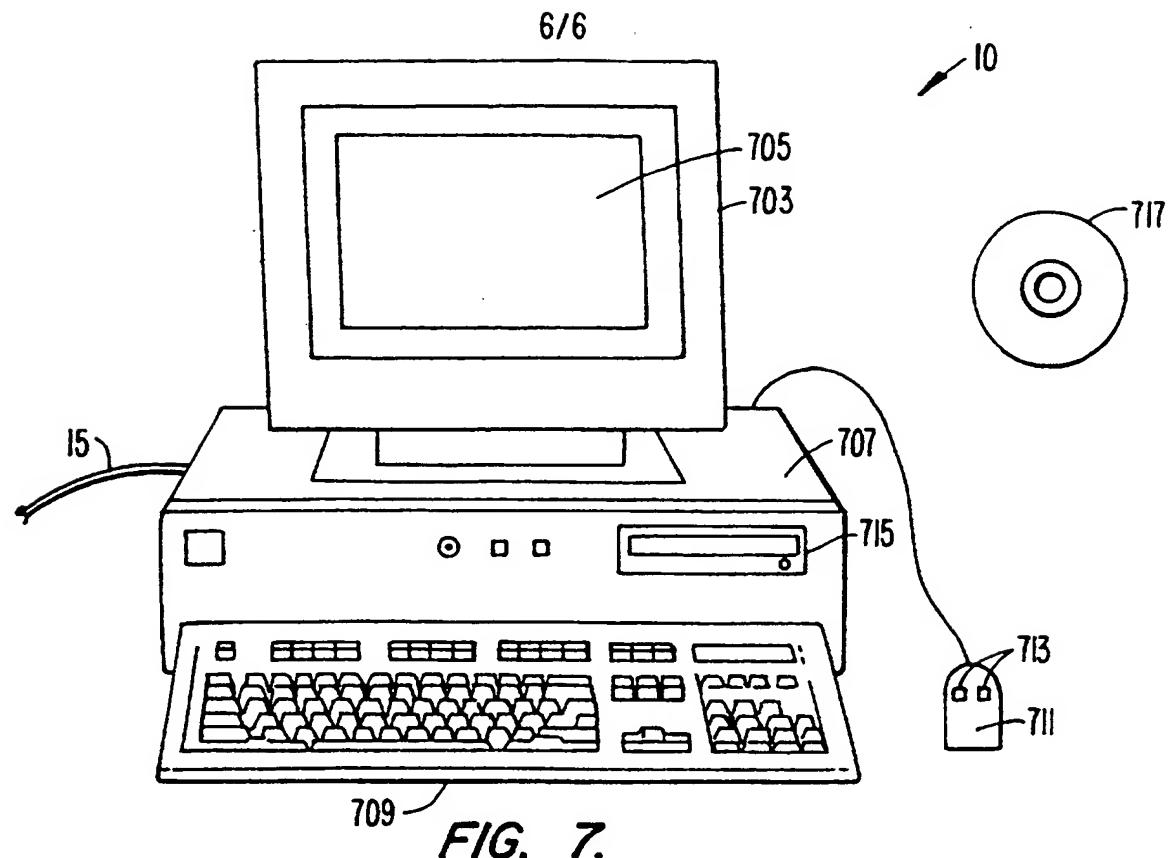


FIG. 6.

**FIG. 8.**

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification <sup>6</sup> :  H04L 29/06		A3	(11) International Publication Number: <b>WO 97/41674</b>  (43) International Publication Date: 6 November 1997 (06.11.97)
(21) International Application Number: PCT/US97/07271  (22) International Filing Date: 29 April 1997 (29.04.97)		(81) Designated States: AU, CA, GB, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(30) Priority Data: 08/641,399 30 April 1996 (30.04.96) US		Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	
(71) Applicant: 3COM CORPORATION [US/US]; 5400 Bayfront Plaza, Santa Clara, CA 95052-8145 (US).		(88) Date of publication of the international search report: 5 February 1998 (05.02.98)	
(72) Inventors: BINDER, James, Stuart; 1277 Elkwood Drive, Milpitas, CA 95035 (US). SIDENBLAD, Paul, James; 10190 Stonydale, Cupertino, CA 95014 (US). BAKER, Roman, Gabriel; 339 Kenmore Avenue, Sunnyvale, CA 94086 (US). CONNERY, Glenn, William; 655 S. Fair Oaks #B301, Sunnyvale, CA 94086 (US).			
(74) Agents: LEBLANC, Stephen, J. et al.; Townsend and Townsend and Crew L.L.P., 8th floor, Two Embarcadero Center, San Francisco, CA 94111 (US).			

(54) Title: PACKET FILTERING BASED ON SOCKET OR APPLICATION IDENTIFICATION

(57) Abstract

A network adaptor driver includes an algorithm for determining special handling of packets by examining socket or port indications for the packets and a means for setting a special handling field in the header of outgoing packets. Specific embodiments include a table for connections that include packets without socket indications, a transmit FIFO list for determining whether a transmit FIFO is full, and a low priority receive queue. A computer readable media allowing a computer to practice the method of the invention is disclosed.

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KC	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

# INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 97/07271

**A. CLASSIFICATION OF SUBJECT MATTER**

IPC 6 H04L29/06

According to International Patent Classification(IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 481 735 A (C.MORTENSEN ET AL) 2 January 1996 see column 2, line 8 - line 34 see column 7, line 53 - column 8, line 24 see column 9, line 41 - column 10, line 22 see figure 6 ---	1,3,5
Y	---	2
X	EP 0 272 834 A (ATT) 29 June 1988 see column 2, line 1 - column 6, line 11 see column 7, line 10 - column 8, line 53 see claim 1 see figures 1,2 ---	10-12
Y	---	14
	-/-	



Further documents are listed in the continuation of box C



Patent family members are listed in annex.

\* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

1

Date of the actual completion of the international search

Date of mailing of the international search report

12 December 1997

19/12/1997

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel +31-701 340-2040, Tx. 31 651 eoo nl.

Authorized officer

Common Anseco

## INTERNATIONAL SEARCH REPORT

Internat'l Application No
PCT/US 97/07271

## C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	WO 94 01828 A (WELLFLEET COMMUNICATIONS) 20 January 1994 see page 2, line 1 - page 3, line 2 see abstract ----	2,14
A	EP 0 669 742 A (HUGHES AIRCRAFT CO.) 30 August 1995 see column 7, line 1 - line 37 ----	1-23
A	D.COFFIELD ET AL: "TUTORIAL GUIDE TO UNIX SOCKETS FOR NETWORK COMMUNICATIONS" COMPUTER COMMUNICATIONS., vol. 10, no. 1, February 1987, GUILDFORD GB, pages 21-29, XP000577221 see the whole document -----	1-23

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

Internat'l Application No

PCT/US 97/07271

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5481735 A	02-01-96	NONE	
EP 272834 A	29-06-88	US 4914653 A CA 1277382 A DE 3751091 D DE 3751091 T JP 1836446 C JP 63234343 A KR 9612686 B	03-04-90 04-12-90 30-03-95 06-07-95 11-04-94 29-09-88 24-09-96
WO 9401828 A	20-01-94	AU 4661793 A US 5463777 A US 5574910 A	31-01-94 31-10-95 12-11-96
EP 669742 A	30-08-95	US 5519640 A	21-05-96